

Bilateral Filter: Graph Spectral Interpretation and Extensions

Akshay Gadde, Sunil K Narang and Antonio Ortega

Ming Hsieh Department of Electrical Engineering

University of Southern California

agadde@usc.edu, kumarsun@usc.edu, ortega@sipi.usc.edu

Abstract

In this paper we study the bilateral filter proposed by Tomasi and Manduchi, as a spectral domain transform defined on a weighted graph. The nodes of this graph represent the pixels in the image and a graph signal defined on the nodes represents the intensity values. Edge weights in the graph correspond to the bilateral filter coefficients and hence are data adaptive. Spectrum of a graph is defined in terms of the eigenvalues and eigenvectors of the graph Laplacian matrix. We use this spectral interpretation to generalize the bilateral filter and propose more flexible and application specific spectral designs of bilateral-like filters. We show that these spectral filters can be implemented with k -iterative bilateral filtering operations and do not require expensive diagonalization of the Laplacian matrix.

Index Terms

Bilateral filter, graph based signal processing, polynomial approximation

I. INTRODUCTION

The bilateral filter (BF) proposed by Tomasi and Manduchi [1] has emerged as a powerful tool for adaptive processing of multidimensional data. Bilateral filtering smooths images while preserving edges, by taking the weighted average of the nearby pixels. The weights depend on both the spatial distance and photometric distance which provides local adaptivity to the given data. The bilateral filter and its variants are widely used in different applications such as denoising, edge preserving multi-scale decomposition, detail enhancement or reduction and segmentation etc. [2]–[6]. Bilateral filtering was developed as an intuitive tool without theoretical justification. Since then connections between the BF and other well known filtering frameworks such as anisotropic diffusion, weighted least squares, Bayesian methods, kernel regression and non-local means have been explored [7]–[11].

This work was supported in part by NSF under grant CCF-1018977.

The BF is data dependent and hence a non-linear and non-shift invariant filter. So, it does not have a spectral interpretation in the traditional frequency domain of images. However we would like to have spectral interpretation so that we can modify the global properties of the signal by changing its frequency components. To overcome this difficulty, we view the BF as a vertex domain transform on a graph with pixels as vertices, intensity values of each node as graph signal and filter coefficients as link weights that capture the similarity between nodes. This graphical views of the BF is used in [6], [9], [12]–[14] but spectral design of filters for images using this graph has not been studied.

We can define spectral filters on these graphs where spectral response is calculated in terms of eigenvectors and eigenvalues of the graph Laplacian matrix [15]–[18]. This spectral interpretation captures the oscillatory behavior of the graph signal [19] and thus allows us to extend of the concept of frequency to irregular domains. This has led to the design of frequency selective filtering operations on graphs similar to that in traditional signal processing. These graph spectral filters also have a vertex domain implementation.

In this paper, we interpret the BF as a 1-hop localized transform on the aforementioned graph. Because the link weights in the graph are data adaptive, the problem of structure preserving filtering boils down to low pass spectral filtering on that graph. We show that the BF can be characterized by a spectral response corresponding to a linear spectral decay. We also calculate the spectral response of iterated BF. We extend this novel insight to build better and more general bilateral-like filters using the machinery to design graph based transforms with desired spectral response. Our design allows one to choose the spectral response of the filter depending on the application which offers more flexibility. We also give a theoretical justification for the design using the framework of regularization on graphs. We show that these spectral filters do not require computationally expensive diagonalization of the graph Laplacian matrix and then provide an efficient algorithm for implementing these filters using the BF as a building block. We examine the performance of the proposed filters in a few applications.

II. BILATERAL FILTER AS A GRAPH BASED TRANSFORM

Consider an input image \mathbf{x}_{in} to the BF. The value at each position in the output image \mathbf{x}_{out} is given by the weighted average of the pixels in \mathbf{x}_{in} .

$$\mathbf{x}_{out}[j] = \sum_i \frac{w_{ij}}{\sum_i w_{ij}} \mathbf{x}_{in}[i] \quad (1)$$

The weights w_{ij} depend on both the euclidean and photometric distance between the pixels $\mathbf{x}_{in}[i]$ and $\mathbf{x}_{in}[j]$. Let p_i denote the position of the pixel i . The weights are then given by

$$w_{ij} = \exp\left(-\frac{\|p_i - p_j\|^2}{2\sigma_d^2}\right) \cdot \exp\left(-\frac{(\mathbf{x}_{in}[i] - \mathbf{x}_{in}[j])^2}{2\sigma_r^2}\right) \quad (2)$$

Spatial Gaussian weighting decreases the influence of distant pixels¹ and intensity Gaussian weighting decreases the influence of pixels with different intensities. Intuition is that only similar nearby pixels should get averaged so that blurring of edges is avoided.

Now, consider an undirected graph $G = (\mathcal{V}, E)$ where the nodes $\mathcal{V} = \{1, 2, \dots, n\}$ are the pixels of the input image and the edges $E = \{(i, j, w_{ij})\}$ capture the similarity between two pixels as given by the BF weights (Figure 1). Image \mathbf{x}_{in} can be considered as a signal defined on this graph $\mathbf{x}_{in} : \mathcal{V} \rightarrow \mathbb{R}$ where the signal value at each node equals the

corresponding pixel intensity. Adjacency matrix \mathbf{W} of this graph is given by $\mathbf{W} = [w_{ij}]_{n \times n}$. Let \mathbf{D} be the diagonal degree matrix where each diagonal element $\mathbf{D}_{jj} = \sum_i w_{ij}$. With this notation the filtering operation in (1) can be written as [9]

$$\mathbf{x}_{out} = \mathbf{D}^{-1} \mathbf{W} \mathbf{x}_{in} \quad (3)$$




Fig. 1: The BF graph

It can be seen from (1) that the output at each node in the graph depends only on the nodes in its 1-hop neighborhood. So, the BF is a 1-hop localized graph based transform. Also, note that the BF includes the current pixel in the weighted average. So, the graph corresponding to the BF has a *self loop* i.e. an edge connecting each node to itself with weight 1. Other filtering techniques such as Gaussian smoothing and non-local means can also be described using similar graphical models [9], [12].

A. Graph Spectrum and Data Adaptivity of the BF

Spectrum of a graph is defined in terms of the eigenvalues and eigenvectors of its Laplacian matrix. The combinatorial Laplacian matrix for the graph G is defined $\mathbf{L} = \mathbf{D} - \mathbf{W}$. We use the normalized form of the Laplacian matrix given as $\mathcal{L} = \mathbf{D}^{-1/2} \mathbf{L} \mathbf{D}^{-1/2}$. \mathcal{L} is a non-negative definite matrix [20]. As a result \mathcal{L} has an orthogonal set of eigenvectors $\mathbf{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ with corresponding eigenvalues $\sigma(G) = \{\lambda_1, \dots, \lambda_n\}$. So, \mathcal{L} can be diagonalized as

$$\mathcal{L} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^t \quad (4)$$

where $\mathbf{\Lambda} = \text{diag}\{\lambda_1, \dots, \lambda_n\}$.

Similar to classical Fourier transform, the eigenvectors and eigenvalues of the Laplacian matrix \mathcal{L} provide a spectral interpretation of the graph signals. The eigenvalues $\{\lambda_1, \dots, \lambda_n\}$ can be treated as graph frequencies, and are always situated in the interval $[0, 2]$ on the real line. The eigenvectors of the Laplacian matrix demonstrate increasing oscillatory behavior as the magnitude of the graph frequency increases [19]. The *Graph Fourier Transform*

¹Eq. (1) shows that the averaging is done over all pixels. However in practice, one assumes non-zero weights only for the pixels which have $\|p_i - p_j\| \leq 2\sigma_s$ [2].




Fig. 2: E_k is the fraction of total energy captured by the first k spectral components.

(GFT) of a signal \mathbf{x} is defined as its projection onto the eigenvectors of the graph, i.e., $\tilde{x}(\lambda_i) = \langle \mathbf{x}, \mathbf{u}_i \rangle$, or in matrix form $\tilde{\mathbf{x}} = \mathbf{U}^t \mathbf{x}$. The inverse GFT is given by $\mathbf{x} = \mathbf{U} \tilde{\mathbf{x}}$.

Figure 2 shows the fraction of total signal energy captured by the first k spectral components of a graph signal corresponding to a 64×64 image block. In this example we consider the following underlying graphs (1) the BF graph (2) the graph corresponding to Gaussian smoothing where the weights depend only on the geometric distance between two pixels. We can see that due to the data adaptivity of the BF graph, most of the signal energy is captured in low frequency part of the BF graph spectrum in comparison with spectrum of Gaussian smoothing graph. So, the spectral basis of the BF graph offers better energy compaction for the given signal.

III. SPECTRAL INTERPRETATION OF THE BILATERAL FILTER

Similar to conventional signal processing, *graph spectral filtering* is defined as

$$\tilde{x}_{out}(\lambda_i) = h(\lambda_i) \tilde{x}_{in}(\lambda_i) \quad (5)$$

$h(\lambda_i)$ is the spectral response of the filter according to which spectral components of an input signal are modulated. Using the definition of GFT and the diagonalized form of \mathcal{L} , we can write graph spectral filtering in matrix notation as [17]

$$\mathbf{x}_{out} = \underbrace{\mathbf{U}}_{\text{Inverse GFT}} \underbrace{h(\mathbf{\Lambda})}_{\text{Spectral response}} \underbrace{\mathbf{U}^t \mathbf{x}_{in}}_{\text{GFT}} = h(\mathcal{L}) \mathbf{x}_{in} \quad (6)$$

To exploit this framework of graph spectral filtering we rewrite the BF in (3) as

$$\begin{aligned} \mathbf{x}_{out} &= \mathbf{D}^{-1/2} \mathbf{D}^{-1/2} \mathbf{W} \mathbf{D}^{-1/2} \mathbf{D}^{1/2} \mathbf{x}_{in} \\ \Rightarrow \mathbf{D}^{1/2} \mathbf{x}_{out} &= (\mathbf{I} - \mathcal{L}) \mathbf{D}^{1/2} \mathbf{x}_{in} \end{aligned} \quad (7)$$

From this equation, we can see that the BF is a graph transform, similar to the one in (6), operating on the normalized input signal $\hat{\mathbf{x}}_{in} = \mathbf{D}^{1/2} \mathbf{x}_{in}$ producing the normalized output $\hat{\mathbf{x}}_{out} = \mathbf{D}^{1/2} \mathbf{x}_{out}$. This normalization allows us to define the BF in terms of the non-negative definite matrix \mathcal{L} and thus have a spectral interpretation. It also ensures that a constant signal when normalized, is an eigenvector of \mathcal{L} associated with zero eigenvalue [21].




Fig. 3: Spectral responses of the BF and iterated BF.

Following (6) we have,

$$\hat{\mathbf{x}}_{out} = \mathbf{U}(\mathbf{I} - \mathbf{\Lambda})\mathbf{U}^t \hat{\mathbf{x}}_{in} \quad (8)$$

This shows that the BF is a frequency selective graph transform with a spectral response $h_{BF}(\lambda_i) = 1 - \lambda_i$ which corresponds to linear decay (See Figure 3). The BF tries to preserve the low frequency components and attenuate the high frequency components.

The BF is used iteratively in many applications. There are two ways to iterate the BF (1) by changing the weights at each iteration using the result of previous iteration (2) by using fixed weights at each iteration as calculated from the initial image. In the first method the BF graph changes in every iteration. So we cannot have a spectral interpretation. In the second method the graph remains fixed at every iteration. Here we consider the second method. Iterating preserves strong edges while removing weaker details. This type of effect is desirable for applications such as stylization [2]. The BF iterated k -times can be written in matrix notation as

$$\mathbf{x}_{out} = (\mathbf{D}^{-1}\mathbf{W})^k \mathbf{x}_{in} = (\mathbf{I} - \mathcal{L}_r)^k \mathbf{x}_{in} \quad (9)$$

where $\mathcal{L}_r = \mathbf{D}^{-1}\mathbf{L}$ is called the random walk Laplacian matrix. It can be shown that any graph transform $h(\mathcal{L}_r)$ can be written in terms of \mathcal{L} as $h(\mathcal{L}_r) = \mathbf{D}^{-1/2}h(\mathcal{L})\mathbf{D}^{1/2}$ [21, Proposition 2]. Using this fact, we can rewrite (9) as

$$\hat{\mathbf{x}}_{out} = \mathbf{U}(\mathbf{I} - \mathbf{\Lambda})^k \mathbf{U}^t \hat{\mathbf{x}}_{in} \quad (10)$$

The spectral responses corresponding $k = 2, 3, 4$ are shown in Figure 3. The figure suggests that iterative application of the BF suppresses more of the high frequency component which is consistent with the observation. Equations (8) and (10) give a different perspective to look at the bilateral filter. They hint at filter designs with better spectral responses that can be tailored to particular applications.

IV. APPLICATION SPECIFIC SPECTRAL DESIGNS

The BF and iterated BF have fixed spectral responses. But these responses may not be suitable for all applications. Below we discuss two applications (1) image denoising and (2) segmentation to illustrate design of more flexible spectral filters.

Denoising.: We consider the problem of image denoising with additive zero-mean white noise.

$$\mathbf{y}[i] = \mathbf{x}[i] + \mathbf{e}[i] \quad (11)$$

where \mathbf{x} is the original image that we want to estimate, \mathbf{y} is the observed noisy image and \mathbf{e} is zero-mean white noise with variance σ^2 . As explained before most of the signal energy lies in the low frequency part of the BF graph spectrum due to its data adaptivity. So most energy in the high frequency spectrum corresponds to noise. We can put this intuition in a more principled framework of regularization where the problem of denoising is equivalent to minimization of a penalty functional [7]. This penalty functional is composed of two terms. The first term is a fit measure and the second term is a data dependent smoothness constraint as captured by regularization operators on the BF graph [22], [23].

$$C(\hat{\mathbf{x}}) = \frac{1}{2} \|\hat{\mathbf{y}} - \hat{\mathbf{x}}\|^2 + \frac{\rho}{2} \|h_p(\mathcal{L})\hat{\mathbf{x}}\|^2 \quad (12)$$

Note that we normalize \mathbf{x} and \mathbf{y} as in (7). The regularization functional can be written in spectral domain as

$$\|h_p(\mathcal{L})\hat{\mathbf{x}}\|^2 = \sum_{i=1}^n [h_p(\lambda_i)]^2 [\tilde{\hat{x}}(\lambda_i)]^2 \quad (13)$$

$h_p(\lambda) \geq 0$ is chosen to be a non-decreasing function in λ so that the high frequency components are penalized more strongly. Putting $\partial C / \partial \hat{\mathbf{x}} = 0$ we get the optimal $\hat{\mathbf{x}}$ as

$$\begin{aligned} \hat{\mathbf{x}}_{opt} &= (\mathbf{I} + \rho h_p^2(\mathcal{L}))^{-1} \hat{\mathbf{y}} \\ &= \mathbf{U}(\mathbf{I} + \rho h_p^2(\mathbf{\Lambda}))^{-1} \mathbf{U}^t \hat{\mathbf{y}} \end{aligned} \quad (14)$$

So, for a chosen regularization functional $h_p(\lambda)$, the spectral response of the denoising filter is given by

$$h_{opt}(\lambda) = \frac{1}{1 + \rho h_p^2(\lambda)} \quad (15)$$

Denoising filters suggested by the regularization framework are essentially low pass filters in the spectral domain of the graph (Figure 5). The BF has a relatively poor spectral decay profile due to which it blurs the textures in the denoising process.

Image Segmentation.: Shi and Malik [24] formulated the problem of image segmentation as a graph partitioning problem on a graph similar to the BF graph. To obtain an m -way partition of the graph, we only need to find the projection of the graph signal on the first few eigenvectors with the smallest eigenvalues. This projection gives a very coarse version of the signal which then can be used to perform image segmentation. So a suitable filter for this application should be a low pass filter in the graph spectral domain with a small cut off frequency and sharp transition band (Figure 6). Iterative application of the BF also gives a coarse version of the image with details

removed. But it favours the eigenvector with the smallest eigenvalue [6] which corresponds to the DC component of the image as seen from its spectral response. On the other hand, for graph partitioning we need the eigenvector with second smallest eigenvalue (and eigenvectors with larger eigenvalues for finer partitions).

To summarize, different applications require filters having different spectral responses $h(\lambda)$. So, we would like to design more general bilateral-like filters with desired spectral response. These filters are of the form $h(\mathcal{L}) = \mathbf{U}h(\mathbf{\Lambda})\mathbf{U}^t$. A direct implementation of these filters requires diagonalization of \mathcal{L} which is of the order $O(N^3)$. For large graphs such as the one considered here, this is computationally very expensive. Fortunately, we can approximate any spectral response $h(\lambda)$ by a polynomial in λ . These polynomial spectral filters are k -hop localized on the graph where k is the degree of the polynomial. This leads to an easy and efficient implementation scheme for these filters as explained in the next section.

V. POLYNOMIAL APPROXIMATION AND FAST IMPLEMENTATION

The spectral response of the iterative bilateral filter (Figure 4(a)) given in (10) is a degree k polynomial. This is a special case of a general class of real polynomials of degree k given as

$$h(\mathbf{\Lambda}) = r_0 \prod_{i=1}^k (\mathbf{I} - r_i \mathbf{\Lambda}), \quad (16)$$

where the roots r_i can be either real or complex conjugate pairs. This generalization allows $k+1$ degrees of freedom in choosing the spectral response of the filter. Further, the corresponding transform \mathbf{H} in pixel domain is a matrix polynomial of \mathcal{L}_r with the same roots as in (16) i.e.,

$$\mathbf{H} = \mathbf{D}^{-1/2} \mathbf{U} h(\mathbf{\Lambda}) \mathbf{U}^t \mathbf{D}^{1/2} = r_0 \prod_{i=1}^k (\mathbf{I} - r_i \mathcal{L}_r). \quad (17)$$

This leads to following result:




Fig. 4: (a) Iterated BF (b) Fast implementation of a bilateral-like filter with polynomial spectral response using BF as a building block

Theorem V.1. *Any graph filter on an image having a polynomial spectral response of degree k can be implemented in the pixel domain as an iterative k step bilateral filter operation.*

Proof: For $k = 1$, the filter $\mathbf{H} = \mathbf{I} - r_1 \mathcal{L} = (1 - r_1)\mathbf{I} + r_1 \mathbf{D}^{-1} \mathbf{W}$. The output in this case is given in (18)

whose convergence rate is r_1 times the convergence rate of the original bilateral filter $\mathbf{D}^{-1}\mathbf{W}$.

$$\mathbf{y} = \mathbf{H}\mathbf{x} = (1 - r_1)\mathbf{x} + r_1\mathbf{D}^{-1}\mathbf{W}\mathbf{x}, \quad (18)$$

For $k > 1$, the transform \mathbf{H} is a cascaded form of k such bilateral filtering operations (See Figure 4(b)). ■

Further, if $h(\lambda)$ is not a polynomial of λ , it can be approximated with a polynomial kernel which can then be implemented as a generalized iterative bilateral filtering operation. It has been shown in [15] that minimax polynomial approximation of any kernel $h(\lambda)$ not only minimizes the Chebychev norm (worst-case norm) of the error between kernel and its approximation, it also minimizes the upper-bound on the error $\|H^{\text{exact}} - H^{\text{approx}}\|$ between exact and approximated filters. In our experiments, we approximate any non-polynomial $h(\lambda)$ with the truncated Chebychev polynomials (which are a good approximation of minimax polynomials).

VI. EXAMPLES

We examine the performance of the proposed spectral design of bilateral-like filters in two applications. First, we consider the image denoising problem. We experiment with a spectral response obtained by the regularization framework in Section 4. We take the regularization functional $h_p(\lambda) = \lambda^2$ which suggests a denoising filter with $h(\lambda) = 1/(1 + \lambda^2)$. We take its 5 degree polynomial approximation. Figure 5 shows the denoising results using this filter and the BF². It can be seen that the BF preserves edges, but it blurs the texture in the denoising process while the proposed denoising filter does a better job of preserving texture. This is also reflected in the SNR values.

Next, we consider an iterative application of the BF. Iterated BF removes minor details from the image while preserving prominent edges. This can be used as an effective preprocessing step in edge-detection and segmentation etc. As stated before, iterated BF favours the DC component which is not useful for segmentation. We use a low pass spectral kernel with small cut-off and sharp transition band so that the second (and a few higher) spectral components get at least as much weight as the DC component. We use a 20 degree polynomial approximation of this kernel and compare it with 20 iterations of the BF. Figure 6 shows that weak edges are blurred more with iterated BF (as expected from the spectral response) compared to the proposed filter.

VII. CONCLUSION

In this paper we explained the bilateral filter as a graph spectral filtering operation. With this novel perspective, we proposed a family of more flexible bilateral-like filters with desired spectral responses. We gave an easy implementation scheme for these filters. Their utility was motivated through few examples. An immediate interesting extension to this work would be to explore different spectral filters suitable for particular applications. Another topic of interest is the design of filter banks using these bilateral-like filters.

²The BF is not the best denoising filter available. We use it in our comparison to emphasize the qualitative differences in filtering results.

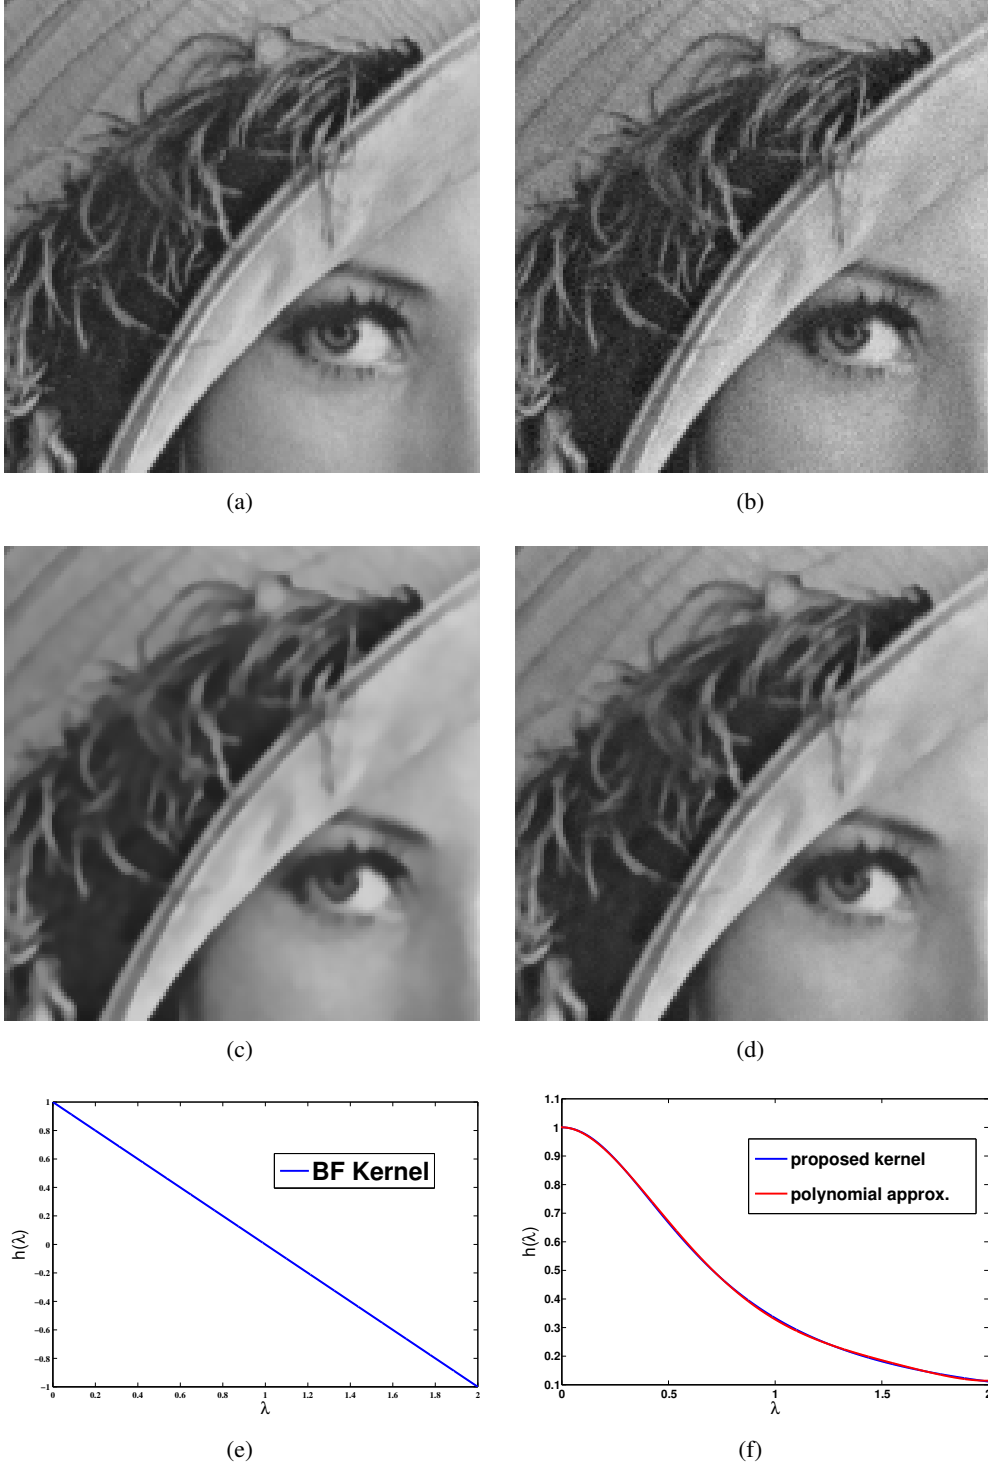


Fig. 5: (a) Original image (b) Noisy image, SNR = 20 dB (c) Output of the BF($\sigma_r = 0.035, \sigma_d = 2$), SNR = 20.65 dB (d) Output of the proposed filter, SNR = 22.64 dB (e) Spectral response of the BF (f) Spectral response of proposed filter

REFERENCES

- [1] C. Tomasi and R. Manduchi, “Bilateral filtering for gray and color images,” in *Sixth International Conference on Computer Vision*, Jan 1998, pp. 839–846.




Fig. 6: (a) Original image (*peppers*) (b) Output of 20 iteration of the BF with changing weights (c) Output of 20 iteration of the BF with fixed weights ($\sigma_r = 0.05, \sigma_d = 2$) (d) output of the proposed spectral filter (e) Iterated BF's spectral response (f) Proposed spectral response

- [2] S. Paris, P. Kornprobst, J. Tumblin, and F. Durand, “Bilateral filtering: Theory and applications,” *Foundations and Trends in Computer Graphics and Vision*, vol. 4, no. 1, 2009.
- [3] Ming Zhang and B. K. Gunturk, “Multiresolution bilateral filtering for image denoising,” *Image Processing, IEEE Transactions on*, vol. 17, no. 12, pp. 2324–2333, Dec 2008.
- [4] R. Fattal, M. Agrawala, and S. Rusinkiewicz, “Multiscale shape and detail enhancement from multi-light image collections,” *ACM Trans. Graph.*, vol. 26, no. 3, 2007.
- [5] F. Durand and J. Dorsey, “Fast bilateral filtering for the display of high-dynamic-range images,” *ACM Trans. Graph.*, vol. 21, no. 3, pp. 257–266, July 2002.
- [6] C. Ye, Y. Lin, M. Song, C. Chen, and D. Jacobs, “Spectral graph cut from a filtering point of view,” *arXiv:1205.4450v2 [cs.CV]*, 2012.
- [7] M. Elad, “On the origin of the bilateral filter and ways to improve it,” *IEEE Transactions on Image Processing*, vol. 11, no. 10, pp.

1141–1151, Oct 2002.

- [8] D. Barash, “Fundamental relationship between bilateral filtering, adaptive smoothing, and the nonlinear diffusion equation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 6, pp. 844–847, 2002.
- [9] P. Milanfar, “A tour of modern image filtering: New insights and methods, both practical and theoretical,” *IEEE Signal Processing Magazine*, vol. 30, no. 1, pp. 106–128, Jan. 2013.
- [10] A. Singer, Y. Shkolnisky, and B. Nadler, “Diffusion interpretation of nonlocal neighborhood filters for signal denoising,” *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 118–139, 2009.
- [11] A. Buades, B. Coll, and J.M. Morel, “A review of image denoising algorithms, with a new one,” *Multiscale Modeling & Simulation*, vol. 4, no. 2, pp. 490–530, 2005.
- [12] G. Peyré, “Image processing with nonlocal spectral bases,” *SIAM Journal on Multiscale Modeling and Simulation*, vol. 7, no. 2, pp. 703–730, 2008.
- [13] G. Noel, K. Djouani, B. Van Wyk, and Y. Hamam, “Bilateral mesh filtering,” *Pattern Recognition Letters*, vol. 33, no. 9, pp. 1101–1107, 2012.
- [14] C. Jian, K. Inoue, K. Hara, and K. Urahama, “Fixed-coefficient iterative bilateral filters for graph-based image processing,” in *Advances in Image and Video Technology*, vol. 5414 of *Lecture Notes in Computer Science*, pp. 473–484. Springer, 2009.
- [15] D. Hammond, P. Vandergheynst, and R. Gribonval, “Wavelets on graphs via spectral graph theory,” *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 129 – 150, 2011.
- [16] S.K. Narang and A. Ortega, “Perfect reconstruction two-channel wavelet filter banks for graph structured data,” *IEEE Transactions on Signal Processing*, vol. 60, no. 6, pp. 2786–2799, June 2012.
- [17] D. Shuman, S. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, “Signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular data domains,” *to appear in IEEE Signal Processing Magazine*, 2013.
- [18] A. Sandryhaila and J. M. F. Moura, “Discrete signal processing on graphs,” *arXiv:1210.4752v2 [cs.SI]*, 2012.
- [19] E.B. Davies, G.M.L. Gladwell, J. Leydold, and P.F. Stadler, “Discrete nodal domain theorems,” *Linear Algebra and its Applications*, vol. 336, no. 1, pp. 51–60, 2001.
- [20] F.R.K. Chung, *Spectral graph theory*, vol. 92, CBMS Regional Conference Series in Mathematics, AMS, 1997.
- [21] S.K. Narang and A. Ortega, “Compact support biorthogonal wavelet filterbanks for arbitrary undirected graphs,” *arXiv:1210.8129 [cs.IT]*, 2012.
- [22] A. Smola and R. Kondor, “Kernels and regularization on graphs,” in *Proceedings of the Annual Conference on Computational Learning Theory*. 2003, Lecture Notes in Computer Science, Springer.
- [23] D. Zhou and B. Schölkopf, “A regularization framework for learning from graph data,” in *ICML Workshop on Statistical Relational Learning*, 2004, pp. 132–137.
- [24] J. Shi and J. Malik, “Normalized cuts and image segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, Aug 2000.